

# Retrofitting Applications with Provenance-Based Security Monitoring \*

Adam Bates  
University of Illinois at  
Urbana-Champaign  
batesa@illinois.edu

Kevin Butler, Alin Dobra,  
Brad Reaves  
University of Florida  
{butler,adobra,reaves}@ufl.edu

Patrick Cable, Thomas  
Moyer, Nabil Schear  
MIT Lincoln Laboratory  
{cable,nabil,tmoyer}@ll.mit.edu

## ABSTRACT

Data provenance is a valuable tool for detecting and preventing cyber attack, providing insight into the nature of suspicious events. For example, an administrator can use provenance to identify the perpetrator of a data leak, track an attacker's actions following an intrusion, or even control the flow of outbound data within an organization. Unfortunately, providing relevant data provenance for complex, heterogeneous software deployments is challenging, requiring both the tedious instrumentation of many application components as well as a unified architecture for aggregating information between components.

In this work, we present a composition of techniques for bringing affordable and holistic provenance capabilities to complex application workflows, with particular consideration for the exemplar domain of web services. We present DAP, a transparent architecture for capturing detailed data provenance for web service components. Our approach leverages a key insight that minimal knowledge of open protocols can be leveraged to extract precise and efficient provenance information by interposing on application components' communications, granting DAP compatibility with existing web services without requiring instrumentation or developer co-operation. We show how our system can be used in real time to monitor system intrusions or detect data exfiltration attacks while imposing less than 5.1 ms end-to-end overhead on web requests. Through the introduction of a garbage collection optimization, DAP is able to monitor system activity without suffering from excessive storage overhead. DAP thus serves not only as a provenance-aware web framework, but as a case study in the non-invasive deployment of provenance capabilities for complex applications workflows.

## 1. INTRODUCTION

Data provenance describes the history of the execution of computing systems, providing detailed explanations as to how data objects were created and came to arrive at their present state. Traditionally, data provenance has been extremely valuable to performing forensics following an attack [5, 24, 41]. For example, provenance can indicate which hosts, processes, files, and data have been affected during the attack and cue cleanup and recovery [37]. Additionally, provenance is of value in virtually any circumstance where a context-sensitive decision must be made about a piece of data. Provenance-aware solutions have been proposed for access controls [31, 32], data leakage [22], malware detection [14], scientific processing [3], and distributed computing [18].

\*This work is sponsored by the Assistant Secretary of Defense for Research & Engineering under Air Force Contract #FA8721-05-C-0002. Opinions, interpretations, conclusions and recommendations are those of the author and are not necessarily endorsed by the United States Government.

Unfortunately, in many complex applications where data provenance would be of greatest value, there is not a general solution for provenance deployment. Although a variety of tools to aid in the design of provenance-aware applications have been proposed [21, 24, 28, 41], modern software is created through the composition of many software artifacts that were written by different developers. Provenance-aware operating systems [7, 15, 16, 29, 34, 35] provide an alternative to ad hoc instrumentation efforts, but are not a complete solution in practice due to semantic gap problems. For example, an operator may want to use provenance to ensure PCI compliance on a credit card database; however, the abstraction level with which the operator wants to work (i.e., credit card records in a database) does not match the system level objects over which provenance is captured (i.e., processes, files, pipes, etc.). This semantic gap also gives rise to the problem of *dependency explosion* – in a long-running program, each output must conservatively be assumed to have derived from all prior inputs [24]. While existing approaches to application layer provenance may overcome one instance of the dependency explosion problem [24, 25], they are not a panacea to this semantic gap problem, as they cannot observe the semantics all components in a complex workflow.

In this work, we introduce a low-cost methodology for retrofitting application workflows with provenance capabilities through a composition of different introspection methods. We present the design and implementation of a unified provenance-aware architecture that includes both novel workflow reconstruction techniques as well as other known approaches to provenance collection. Our exemplar provenance-aware workflow mechanism, DAP,<sup>1</sup> is designed with consideration for the unique challenges and opportunities presented by web service environments. DAP is a transparent collection agent that captures detailed provenance of application workflows without suffering from the semantic gap problems of system-level collection. Our approach leverages minimal knowledge about common workflow structures in order to extract precise and efficient provenance. In particular, DAP leverages the widespread adoption of the SQL syntax, transparently interposing on database connections to interpret and extract the provenance of database transactions. DAP is compatible with a large percentage of existing web services, generating concise and understandable provenance with little or no system modification.

Our contributions can be summarized as follows:

- We present the design and implementation of DAP, a minimally invasive, low overhead framework for capturing workflow provenance. DAP combines state-of-the-art provenance

<sup>1</sup>Dapping is a form of fly fishing that causes minimal disturbance to the water. Likewise, our DAP system is minimally invasive to application workflows while extracting precise contextual metadata.




Figure 1: A provenance graph showing the actions of a potentially malicious binary that is running with root privileges. Edges encode relationships that flow backwards into the history of system execution, and writing to an object creates a second node with an incremented version number.

techniques with novel components that aggregate the provenance of application objects. We also address the challenging problem of integrating provenance from different capture points under a common namespace.

- We present an extended case study through which we demonstrate DAP to be an effective means of reasoning about, detecting, and actively preventing Internet-based attacks. In particular, we demonstrate that our system can be used as a means of preventing SQL Injection (SQLi)-based data exfiltration, one of the most widespread and insidious threats to the Internet today. We also show how DAP can be used in concert with other technologies to track system layer attacks against the web server.
- In evaluation, we show that our implementation imposes just 5.1 ms of overhead on web application requests, and microbenchmark individual steps in our system to arrive at a better understanding of this cost. Automatic provenance collection is known to impose excessive storage overheads on long running systems; however, through applying a garbage collection optimization, we show that our mechanism can monitor for active SQLi attempts while maintaining sublinear growth in storage burden.

## 2. WEB APPLICATION PROVENANCE

Data provenance describes the actions taken on a data object from genesis onwards, including how it came to exist in its present state. Provenance can be queried to answer questions such as “*What datasets were used in the creation of this data object?*” and “*In what environment was this data object produced?*” A standard representation for data provenance is a directed acyclic graph which is specified in the W3C PROV data model [39], which we will use throughout this work. An example provenance graph plotting the execution of a potentially malicious binary is shown in Figure 1. This binary, while running with root privileges, first read several system files, including `/etc/shadow` and `/etc/rc.local`. It then wrote to those files in an attempt to gain persistent access to the system. In the graph, edges represent relationships between different system objects. To prevent cycles from forming in the graph, writing to an object triggers the creation of a new node that represents a new version of the object.

Web applications present a challenging scenario for provenance because of their heterogenous nature; web requests traverse the operating system, web server, web application, and database management system, each of which maintain their own internal semantics. Consider the representative scenario in Figure 2. There are  $i$




Figure 2: Diagram of a web service architecture. To accurately track provenance, it is necessary to track individual client requests from the network, through the server, to the database, and back.

different Internet clients sending requests. Network requests pass from the the operating system to the web server software, which the server handles concurrently with  $j$  different workers. The workers’ database transactions are then multiplexed between  $k$  different connections. Unfortunately, there is not necessarily any equivalence between the numbers  $i$ ,  $j$ , and  $k$ .

To reason about the attack service of a web application it is necessary to understand the actions of each of these components, yet deploying provenance capabilities in this domain is particularly challenging due to the complexity of this workflow. Provenance-aware operating systems such as PASS [30], LPM [7], and ProTracer [27] provide a single point of observation for all system activity, but are not a complete solution due to the semantic gap that divides the system and application layers. For instance, these systems would struggle to disambiguate web server requests as autonomous units of work (i.e., *dependency explosion* [24]), leading to the false conclusion that each server response was dependent on all previous client requests. Lee et al.’s LogGC [25] and BEEP [24] system provide space-efficient forensics for application monitoring, but suffer from the same semantic gap problem as provenance-aware operating systems due to their reliance on system audit logs.

The remaining alternative to the above approaches is to undertake a tedious instrumentation effort of the web application. Although provenance libraries exist that simplify the manual instrumentation of source code [28], this approach requires additional resources and domain-specific knowledge that is unlikely to be available to most web developers. Furthermore, instrumentation could extend past the primary software artifact to its dependencies, including the web server, runtime framework, and other third party libraries. This solution may even require re-architecting the web service use a provenance-aware database management services such as Trio [38], DBNotes [9], and ORCHESTRA [23]. Due to the extraordinary capital required by this approach, we conclude that is not a viable solution to creating provenance-aware web services. What is needed instead is a means of retrofitting provenance into existing services with minimal cost to web developers.

## 3. DESIGN

### 3.1 Threat Model & Assumptions

The attack surface we consider in this work is that of a typical web application. By connecting to the application’s external listening ports the attacker may attempt a variety of misdeeds on the system. The attacker may attempt to exfiltrate data from the web application through iterative command injection (e.g., SQLi)




Figure 3: The provenance graph of a typical web session that we wish to observe with our architecture.

attacks. A successful exfiltration attack will involve repeated command injections as the attacker attempts to discover the location of valuable data. The attacker may also attempt to use command injection to inject spurious data into the database for the purposes of privilege escalation or cross-site scripting. Alternately, the adversary’s target may not be the web application but the system itself. The attacker may be attacking the web server in order to compromise other services on the host or to move laterally to other hosts on the network [40].

We make the following assumptions about the security of each web service component. We conservatively assume that the web server, web application, and database engine are all subject to compromise. These components may begin to lie about their actions on the system at any time, but we assume that at least one provenance record of the attacker’s access attempt is recorded prior to compromise. We also assume that the integrity of the host kernel is assured. This condition is made more reasonable through the deployment of kernel hardening techniques, integrity measurement, and mandatory access control (e.g., SELinux) that protects the operating system’s trusted computing base. Finally, we assume that the novel components our system introduces are *not* subject to compromise. As we will later show, these mechanisms are small and simple enough to be subjected to rigorous audit, and can also be protected through system hardening techniques.

## 3.2 System Goals

- G1 Complete.** Our system must offer a complete description of individual requests as they pass through an application workflow. The record must remain complete in the presence of unexpected events triggered by attacker behavior, such as command injection attacks or binary exploitation. If we elect to forego provenance capture at a given system component, the recorded provenance must provide sufficient context to reconstitute the entire workflow.
- G2 Integrated.** Our system must combine provenance from different operational layers in a salient manner that provides a coherent explanation of application activity to the administrator. Provenance generated by different capture agents must share a common namespace, and each capture agent must be able to accurately reference the activities of other agents.
- G3 Minimally Invasive.** Provenance, like security, is often perceived as a cost burden. Our system must therefore impose a bare minimum number of modifications to existing system components, including the application and backend infrastructure (e.g. database engine, web server). Optimally, our solution would not make any changes to existing software, instead introducing independent mechanisms so that the system would continue to function correctly as software in the application is periodically upgraded.
- G4 Widely Applicable.** To further advocate for the deployability of provenance-aware applications, our efforts in the development of the system should not be limited to the benefit

of a particular application, backend component, or architecture. Instead, our system should be immediately compatible with a broad number of existing applications.

- G5 Defensive Capabilities.** While provenance is invaluable to forensic investigation after an attack has occurred, attacks on Internet domains are frequent and relentless. Therefore, our system must be fast enough to provide real-time assistance to the defense of the host. This includes the ability to detect and explain attacks as they occur and aid in system recovery in the event of a successful attack.

### 3.2.1 Provenance Definition

As provenance is codified in dramatically different ways throughout the literature, from exhaustive descriptions of system activity [34] to lightweight proofs of program execution [26], an important step in the design of our system is to identify the scope and granularity of the events we wish to observe. As the ultimate goal of our system is to observe the attacker described in Section 3.1, the provenance we collect must exhaustively describe the manner in which attacker inputs interact with the application workflow. We illustrate this with a typical web application, where we must be able to track a client request from receipt on the host, through a specific worker in the web server, through a database request and response, until a response is crafted by the web application and returned to the remote client. We must also be able to differentiate between different client requests, even if they are performed by the same worker or re-use the same database connection.

In order to satisfy Goals G3 and G4, our system must avoid modifying the database management system. As a consequence of this, we will be unable to know the precise records that are impacted by a query. Therefore, we must describe SQL objects not at the granularity of database records, but instead as database columns and tables, which can be inferred from the query itself. As we will show in Section 3.6, this coarser granularity is well suited to explaining command injection techniques, which often involve access to columns that should never be returned to the user.

With this in mind, the goal of DAP is to produce provenance explanations for individual web requests like the one shown in Figure 3. Network activity is tracked at the system granularity, the web application is tracked at the granularity of individual units of work, and database objects are tracked as columns and tables. Provenance captured at different sources will be integrated through their shared relation to the web application worker during a given unit of work.

## 3.3 Provenance Capture: Overview

In order to achieve the above goals, what is needed is a comprehensive provenance architecture that is able to understand the semantics of the web server, database, and operating system in unison. An overview of our solution to this challenge, DAP, is shown in Figure 4. Provenance-aware components are shaded in orange. DAP introduces several provenance-aware components, but requires no modification to the Web Application or Database Engine. The provenance-aware components are a small and re-usable modification to the Web Server to facilitate execution partitioning,

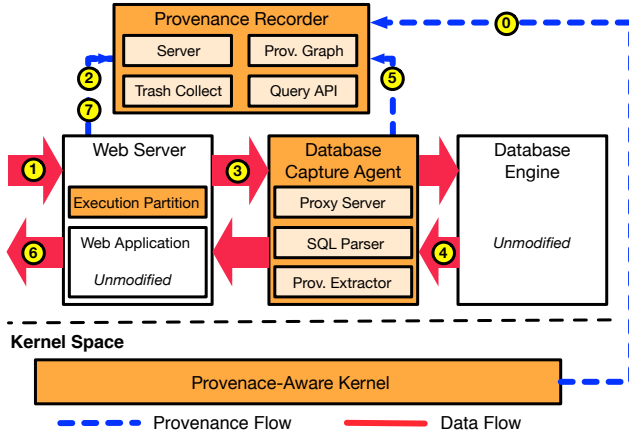


Figure 4: Overview of the DAP architecture. Provenance-Aware components are shaded in orange. No changes are required to the Database Engine or Web Application; instead, provenance is generated by interposing on the connection between the Web Application and Database Engine. A small change to the Web Server is required to facilitate execution partitioning.

a Database Capture Agent that transparently proxies all traffic between the Web Application and the Database Engine, and a Provenance Recorder that aggregates provenance information between the different parties. Our system also assumes that system layer provenance is being collected, which can be obtained through use of a custom provenance-aware kernel [7, 29, 34, 35] or user-space system monitor [15, 16].

The workflow for provenance collection is as follows: (1) a remote host makes a request to the Web Application; (2) a small modification to the Web Server performs *execution partitioning* [24], notifying the Provenance Recorder whenever the Web Application has started a new autonomous unit of work; (3) the Database Capture Agent proxies and subsequently parses a query issued to the Database Engine; (4) after measuring the impact of the query by parsing the Database Engine response, the Database Capture Agent (5) transmits provenance information to the Provenance Recorder; (6) as the Web Applications transmits a response to the remote host, (7) the Web Server notifies the Provenance Recorder that the unit of work has ended; throughout execution, (0) the provenance-aware kernel generates provenance for all activities that are not being explicitly disclosed by the Web Server or Database Capture Agent.

In the remainder of this section, we will describe in greater depth the operation of the Database Capture Agent as well as introduce a provenance-based defensive mechanism. The Execution Partition and Provenance-Aware Kernel components rely on known techniques and are discussed at greater length in Section 4.

### 3.4 Provenance Capture: Database

A fundamental design consideration in our system was the manner in which DAP would observe communication between the Web Application and Database Engine. One possibility would be to instrument the web service to extract database queries. This would have made our solution application-specific, violating Goal G4. Another possibility would be to instrument the database, or to use an existing provenance-aware database. However, instrumenting a database engine would also limit our solution to the benefit of a particular database service, violating Goal G4. In turn, using a provenance-aware database would require re-architecting the web

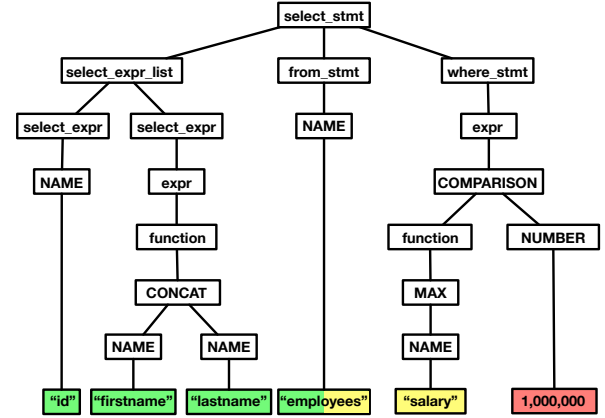


Figure 5: A simplified example of a SQL parse tree for the statement “*SELECT employee\_id, CONCAT(firstname, lastname) FROM employees WHERE MAX(salary) > 1,000,000*”. The leaves of the tree are color-coded by their provenance extraction decision condition. Columns in the select expression are noted by a `SQL_READ` provenance event. Columns present in other subexpressions are noted by a `SQL_USED` event. Non-persistent entities such as numbers, and functions are not considered for provenance extraction.

service, violating Goal G3.

Instead, we chose to implement an explicit TCP proxy that interposes on communications between the application and database. The only change required by this approach is that either the application or the database change the port over which they communicate with one another, allowing the proxy to interpose. These types of configuration options are nearly always exposed and easily modifiable in both database and web application software. We rejected a fully transparent solution that used `iptables` to capture packets between the components, as it would substantially increase the complexity of the capture agent. We also rejected achieving interposition by modifying an existing database connection library as this would limit the applicability of our agent to applications with that particular dependency, violating Goal G4.

#### 3.4.1 Query Parsing

After proxying the web application’s traffic, we make use of a query grammar and parser to extract fine-grained provenance information from database queries. We chose to focus on the SQL language in this work due to its widespread use. While in reality SQL comes in many different flavors and varies by database management system, our needs are foundational enough that the syntactical differences between SQL variants can be largely ignored.

The output of a SQL parser can be visualized as a parse tree, a simplified example of which is shown in Figure 5. This tree is a `SELECT` statement that contains a `FROM` clause (required) and a `WHERE` clause, which is one of several optional clauses. We use this example to demonstrate how DAP handles the various data objects contained in a SQL query:

**Data Accessed:** We refer to the named objects referenced in the primary clause of the query as accessed data. These are the objects that will be returned by the database in its response to the query. When a query is parsed, DAP generates a `SQL_READ` provenance event for each piece of accessed data. The accessed data in Figure



5 are the *employee\_id*, *firstname*, and *lastname* columns, and the *employees* table.

**Data Referenced:** Named objects that appear in subsequent clauses of the query are not explicitly returned by the database, but nonetheless inform the response message. Consider again the example query in Figure 5 – while employee salaries are not returned in the query, the response implicitly informs the querier of which employees salaries are greater than \$1,000,000. Referenced data therefore represents a dangerous side channel for information leakage. However, in some environments it may be unnecessarily conservative to treat all referenced data as accessed data. To account for this, we introduce a `SQL_USED` event to describe referenced data.

**Ephemeral Data:** Query expressions also include non-persistent data objects, such as numbers and string literals. In the case of `SELECT` statements, ephemeral data can manipulate the records and values returned by a query, but not the columns accessed and referenced. We choose to ignore ephemeral data for the case of provenance extraction.

During parsing, the SQL grammar tracks named and referenced data via synthesized attributes. The name, and also the prefix if present, is added to a linked list as the statement is parsed. At the root of the statement, a function determines the appropriate prefix for each column given the tables used in the `FROM` clause.

While we use the `SELECT` statement in the scenario above, the same rules can be applied to other statements. `SHOW` and `DESCRIBE` statements can be treated the same way as `SELECT` statements. For expressions that write to the database, we introduce the provenance event `SQL_WASGENERATEBY`. This event is used to describe the column and table references that appear in the primary clauses of `INSERT` and `UPDATE` expressions. This event contains the same fields as the `SQL_READ` event, but is handled differently by the Provenance Recorder. When a `SQL_WASGENERATEBY` is received, the Recorder will create a new node for the accessed object with an incremented version number. Subsequent `SQL_READ` and `SQL_USED` events will be linked to the newer node in order to prevent cycles from forming in the graph. While we do not explicitly address any other statement types in this work, our rules generalize to any expression that reads from or writes to the database.

### 3.4.2 Parsing Challenges

When extracting provenance from non-trivial SQL statements, a variety of challenges arise. We came across a number of such challenges while designing and implementing DAP. We describe our solutions to each problem below:

**Parsing Challenge #1: Wildcards.** Through use of the wildcard character, SQL statements are able to reference all columns in a table without explicitly naming them. To address this, we provide the Provenance Recorder a schema description, which allows it to translate the wildcard character into the associated columns for the given table. In our implementation, we obtain the schema through use of the `mysqldump` command.

**Parsing Challenge #2: Aliases.** Any value in a SQL statement can be aliased to another name. The challenge in resolving aliases is that an alias may be referenced in one clause of the query, but defined in another. To address this problem, our SQL grammar makes use of synthesized attributes to track references and definitions of aliases. At the top level of the parse tree, the list of referenced aliases are then resolved to their true table and column names. In effect, this means that DAP unaliases named objects during parsing, ensuring that the extracted provenance is unobfuscated.

**Parsing Challenge #3: Nested Queries.** An additional obstacle we faced in the design of our grammar was that of nested queries. In SQL, full statements can be indefinitely nested within one an-

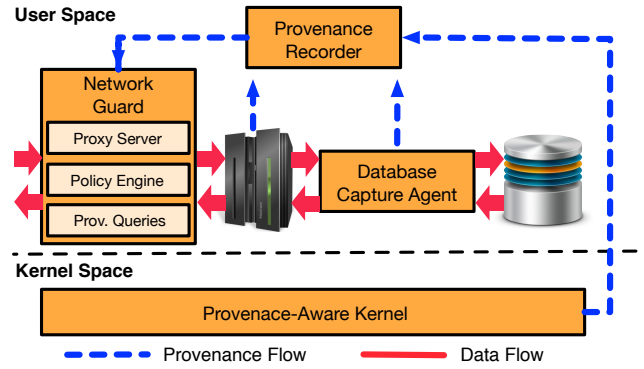


Figure 6: We extend DAP with analysis components that provide real-time detection of data exfiltration attacks (Network Guard).

other. For example, “`SELECT A FROM (SELECT id AS A FROM employees)`” is a valid statement. Nested queries can be used to further obfuscate the true origin of a data object. Our solution to this is to modify the synthesized attribute routines described above. At the root of each subquery in the parse tree, objects in the subquery are unaliased, and the named and referenced objects used by the subquery are transferred to the parent query. Additionally, the alias mapping is passed to the parent query. This allows DAP to unnest queries as the statement is parsed.

## 3.5 Provenance-Based Security Mechanism

Given the above provenance capture agents, DAP can be used to thwart SQLi-based data exfiltration attacks in real-time by performing provenance-based whitelisting of outbound server responses. We accomplish this through introducing a Network Guard component to the DAP architecture as shown in Figure 6. The Network Guard is another TCP proxy server placed between web server and the network. It takes no action on incoming traffic, but inspects outbound network traffic from the server. When the guard intercepts a transmission to the remote host, it examines the network context to identify the process id of the worker thread. It then issues a query to the Provenance Recorder to obtain the list of data ancestors used by the worker during the current unit of work.

In order for the Network Guard to prevent exfiltration, we require a means of encoding developer or administrator intent about web service workflows. To do so, we introduce a simple policy language. The Network Guard accepts a list of policy rules and checks the list of ancestors against each rule before permitting the message to be transmitted. Rules take the form:

$$< RULE, [Tab_1.Col_1, \dots, Tab_n.Col_n], SIZE >$$

*RULE* can be set to *ALLOW* to whitelist ancestors, or *DENY* to blacklist ancestors. A list of SQL objects can then be specified to match against in the response messages ancestry. We support lists of objects, rather requiring multiple policy entries for each object in the list, because at times the fusion of data objects reaches a different sensitivity level than the level of each individual object (e.g., Personally Identifiable Information<sup>2</sup>). *SIZE*, expressed in bytes, approximates the impact of a query. It can be used to limit the amount of SQL data that should be returned to a client in a response message, or can be disabled by setting to 0.

## 3.6 Security Analysis

<sup>2</sup> See NIST SP 800-122

Our arguments for the security of DAP are as follows:

**Complete (G1).** DAP is able to differentiate between individual web requests through a minimal modification to the Apache 2 server that performs execution partitioning. Transactions with the database are observed by DAP's Database Capture Agent. Because the port for the database engine is hard-coded into the web application, there is no other means to reach the database then to go through the Capture Agent. In the event that the server is compromised, DAP can continue to track the actions of the attacker on the system through use of the provenance-aware kernel, whose trusted computing base can be isolated from the rest of user space using SELinux enforcement [7].

**Integrated (G2).** A major challenge in layered provenance systems is establishing a uniform namespace for objects [2]. In DAP, we integrate provenance from different capture agents through tracking process ids (pid). In our implementation, the web server is configured to run in pre-fork mode in which all workers receive a unique pid. When a request is proxied by the Database Capture Agent, the worker's pid is recovered through matching the network context in the `sockaddr_un` struct against the list of active tcp sockets returned by `netstat`. The pid is then embedded in all subsequent provenance events generated by the Database Capture Agent. To integrate DAP provenance with system-layer provenance, we introduce a `pid_to_provenance` syscall that accepts a pid and returns the universally unique identifier associated with the provenance of the process' fork in the kernel. Thus, all layers share a common language to describe an activity.

**Minimally Invasive (G3).** DAP works without requiring any changes to existing web applications. In our implementation, the web server is minimally modified through the introduction of a new header file as well as the insertion of 3 lines of code in the existing source. The database requires no change except to be reconfigured to listen on another port, which can be done without recompilation. If the ability to track attacker actions after a server compromise is desired, a provenance-aware kernel must be installed on the machine. However, DAP is able to track service layer attacks, such as command injection, without this capability.

**Widely Applicable (G4).** We confirmed that our DAP implementation is compatible with the Apache 2 and Tomcat (via `mod_jk`) web frameworks, as well as with the MySQL and PostgreSQL database engines. DAP's use of the SQL grammar is simple enough that it should work with any SQL variant with only minimal modification. In Section 7, we discuss whether our approach generalizes to NoSQL databases.

**Defensive Capabilities (G5).** The Network Guard component proxies outbound network traffic and searches its provenance for evidence of data exfiltration. An administrator can use this tool to prevent certain data objects from ever being returned to the client, or even to terminate connections if a suspicious amount of data is being transmitted to the client. Moreover, we reduce the attack surface of DAP by isolating the complex parsing procedure from its parsing responsibilities, making it more difficult to disable.

## 4. IMPLEMENTATION

We have implemented DAP for the Linux operating system. We have tested that our system works with both MySQL and PostgreSQL by using their command line clients. In both cases, the only change required to the database engine was to modify the port on which they listen for connections. We also confirmed that DAP worked correctly with a variety of web applications and tools

including MediaWiki<sup>3</sup> (PHP-based), UnixODBC<sup>4</sup>, and a suite of Tomcat-based applications released with the AMNESIA [19] evaluation testbed<sup>5</sup>.

### 4.1 Provenance-Enhanced Web Server

We have instrumented the Apache 2 web server (version 2.2.31) to perform execution partitioning on each TCP socket. The server is configured to run in pre-forked mode<sup>6</sup>. Because we did not have access to the BEEP tool, we manually instrumented the source code. To minimize the impact on the rest of the Apache 2 code base, all of the logic required to report execution partitions to the Provenance Recorder was included in a single header file. As a result, we inserted just 3 lines of code into the existing source files. The changes were made to `server/config.c`. The first insertion included our header file. The second two lines are placed before and after the call to `ap_run_handler` in the `ap_invoke_handler` function:

```
/* DAP -- Transmit "Unit Start" Message here! */
char * uuid = dap_unit_start(
    r->connection->remote_addr);

/* Handle the request */
result = ap_run_handler(r);

/* DAP -- Transmit "Unit End" Message here! */
dap_unit_end(uuid);
```

The `dap_unit_start` function generates a UUID to associate with the unit of work, then transmits the `UNIT_START` message to the Provenance Recorder that contains the UUID and the `remote_addr` struct. The `dap_unit_end` function transmits a `UNIT_END` message to the Provenance Recorder that contains the UUID, then frees the UUID character array.

Instrumenting this layer of the Apache 2 stack offers several advantages. First, it resides beneath the various error and security filters performed by the server, ensuring that we do not generate provenance for requests that are later rejected. Second, it resides above the file-specific handler module, so we are able to instrument both static web pages and dynamic web applications. As a result, our provenance-enhanced Apache 2 works for various handler modules, not just HTTP. For instance, when we later recompiled Apache with the `mod_jk` module<sup>7</sup>, we found that our code also worked on Tomcat applications without any modification.

### 4.2 Database Capture Agent

Our capture agent is a multithreaded TCP proxy server that listens on the database engine's assigned port. Once connected, the server extracts database queries issued by the web application. It then passes them through a Bison parser that makes use of a publicly available SQL grammar<sup>8</sup>. We extended the grammar to aggregate the columns and tables accessed by the query as described in Section 3.4.1.

When a new connection is proxied, the Capture Agent first recovers the process ID (pid) of the sender by matching the network context in the `sockaddr_un` struct against the list of active tcp sockets returned by `netstat`. After parsing the query, the Capture Agent inspects the list of database objects accessed. It then

<sup>3</sup> Available at <https://www.mediawiki.org>.

<sup>4</sup> Available at <http://www.unixodbc.org>.

<sup>5</sup> Available at <http://www-bcf.usc.edu/~halfond/testbed.html>.

<sup>6</sup> See <http://httpd.apache.org/docs/2.2/mod/prefork.html>.

<sup>7</sup> See <http://tomcat.apache.org/connectors-doc>.

<sup>8</sup> Available at <https://github.com/hoterran/sqlparser>.

creates a provenance event for each object, which is a tuple of the form  $\langle pid, relationship, column, table \rangle$ . *Relationship* is one of the several relationships specified in Section 3.4, and are also consistent with the W3C PROV-DM model. Once the new provenance event is created, it is sent to the recorder through a Unix socket. Rather than design our own protocol, we extend the space-efficient Hi-Fi protocol [34] to support several new kinds of events.

In order to ensure application stability, the Database Capture Agent’s parsing functionality is isolated from the TCP proxy server. The parser is implemented as a separate binary that is invoked by the proxy using the `system` syscall. After running the parser, the proxy checks its exit code status. If the parser exited in a bad state, indicating a potential attack, the proxy drops the connection and transmits the input to the Provenance Recorder as an attribute to be added to the provenance graph. Because the proxy process does not directly inspect messages from the application, we are confident that it cannot be disabled by malformed inputs.

### 4.3 Provenance-Aware Kernel

In the event that the Web Server is compromised, the above components alone are insufficient to track attacker actions. This is because the attacker will no longer be limited to the Web Application workflow, but will instead be able to take any action on the system with the privileges of the Web Server. More dangerously, the attacker may be able to use this foothold to escalate to root privileges. Notably, all application-based provenance-tracking [24] suffers from the same limitation of being unable to reliably record provenance after the point of compromise. To ensure the ability to track attacker actions after a server exploit, we ran our application on top of a provenance-aware operating system. We made use of the Linux Provenance Modules project for our provenance-aware kernel. We configured LPM to make use of the Hi-Fi module [34], and deployed the system as described in [7].

### 4.4 Provenance Recorder

The Provenance Recorder is responsible for aggregating provenance from the different capture agents and representing it in an efficiently queried in-memory graph. We implemented the Recorder in C++ using the SNAP graph library. The Recorder listened for new provenance events over Unix sockets. Different provenance events are handled as described in Section 3.4; generally speaking, when the Recorder received a new event it first checked to see if any of the involved objects were already present in the graph, created them if they are not, and then added a new relationship between the objects. Visual examples of how provenance graphs were represented by the recorder follow in Section 6.

### 4.5 Network Guard

Like the Database Capture Agent, the Network Guard is a multithreaded TCP proxy server that is placed between the web server and the network. When outbound traffic is transmitted from the web server, the Network Guard issues an ancestry query request to the Provenance Recorder by using the pid of the sending worker as a unique identifier. Upon receipt of the list of ancestors, the Network Guard verifies policy compliance prior to permitting the data to be transmitted over the network.

## 5. EVALUATION

We evaluated DAP using a VMWare Fusion VM running CentOS 6.5. We executed our webserver tests in the common deployment model of a virtual machine with 4GB RAM and 2 vCPUs. The host was a local server with two 2.4 GHz Quad-Core Intel Xeon processors and 12 GB RAM.

Scenario	Time
Web Application w/o DAP	34.5 ms
Web Application w/ DAP	29.3 ms
Overhead	5.1 ms (17.1%)

Table 1: End-to-end delay imposed by DAP during the DVDStore Benchmark

Location	Operation	Time
Apache 2 <code>config.c</code>	Transmit <code>unit_start</code> message	0.82 ms
Apache 2 <code>config.c</code>	Transmit <code>unit_end</code> message	0.91 ms
Database Capture Agent	Parse SQL Query	0.11 ms
Database Capture Agent	Transmit SQL Provenance	1.98 ms
Database Capture Agent	Other (incl. proxy cost)	1.28 ms

Table 2: Microbenchmarks for DAP system during the DVDStore benchmark

### 5.1 End-to-End Delay

One of the vital measures of DAP’s performance is the end to end delay it imposes on web requests. The apparatus we used for both this test and its subsequent microbenchmarks was the Dell DVD Store Database Test Suite,<sup>9</sup> an open source simulation of an online ecommerce site. We configured DVDStore to run on MySQL with a 10 GB database. DVDStore’s benchmarks are issued directly to the DBMS, bypassing the web front-end of the ecommerce site; in order to measure our modifications to Apache 2, we modified DVDStore’s client workload driver to issue all SQL queries as curl requests to port 80 of the host. The queries were received by a toy PHP application on Apache 2 that relayed SQL queries to MySQL using POST requests, then returned the results to the client. In total, the DVDStore workload issued over 10,000 unique SQL queries to our system.

We measured overall performance under two configurations. In the first (without DAP), an unmodified copy of `httpd` communicated directly with MySQL. In the second (with DAP), our modified `httpd` communicated through the database capture agent. Table 1 summarizes our findings. The average response time for queries when DAP was disabled was 29.3 ms. The average response time for queries when DAP was enabled was 34.5 ms, representing a cost of just 5.1 ms, or 17% overhead.

### 5.2 Microbenchmarks

During the above trial, we instrumented DAP to measure the time spent on individual steps involved in provenance capture – the `unit_start` and `unit_stop` messages generated by the web server, the SQL parsing step, and the transmission of provenance from the Database Capture Agent to the Provenance Recorder. By subtracting these measures from an end-to-end measurement, we also captured the approximate cost of other steps, most notably the delay imposed by proxying database traffic. The results are shown in Table 2, and Figure 7 shows the associated cumulative density functions for each measure. The various SQL queries generated by DVDStore could be parsed and have their provenance extracted in an average of 0.11 ms. The primary source of delay in our DAP system is due to inter-process communication; transmitting small provenance events to the recorder required approximately 1 ms, and larger messages required approximately 2 ms. Steps that required inter-process communication experienced high variance, indicating processing delays at the Recorder that could be addressed to improve performance. As our provenance recorder implementation was single-threaded, it is likely that delays could be dramatically

<sup>9</sup>See <http://linux.dell.com/dvdstore/>




Figure 7: Cumulative Densities of each microbenchmark step during the DVDStore benchmark.




Figure 8: Cumulative Density of Provenance Recorder's response time when queried by the Network Guard.

reduced through creating a multi-threaded version of the recorder.

### 5.3 Network Guard Performance

Benchmarking our Network Guard involved determining the speed with which DAP provenance can be analyzed in a live system. This was the first evaluation trial in which the Network Guard was enabled. We instrumented the Network Guard to measure the time required to query the Provenance Recorder and process its response message. For this trial we repeated the DVDStore benchmark, which included a great diversity of SQL queries. The results are shown in Figure 8. The average response time by the Provenance Recorder was just 1.23 ms. We microbenchmarked this result as well, and found that on average 1.17 ms of this delay was due to IPC, while just 0.5 ms was required to generate the provenance ancestry. In the worst case, the query took 7 ms to respond, but this was also due to IPC delays and not to the cost of graph traversal. These results indicate that even our proof-of-concept implementation can be used as an enforcement mechanism without imposing unacceptable latency.

### 5.4 Storage Overhead & Optimization

A vital consideration when collecting data provenance is the storage overhead incurred. Not only do high overheads increase the cost of storage, but preserving unnecessary provenance impacts the speed with which the provenance can be queried. To capture the storage overhead of DAP, we observed the growth of the provenance graph during the DVDStore workload. Every 500 ms, we




Figure 9: Growth of provenance storage during the DVDStore benchmark

polled the total size of the provenance recorder's memory allocation using the `/proc` file system. The results of this trial are shown in Figure 9. During the first 6 minutes of the DVDStore trial, the provenance logs grew to 15.7 MB in size. Unfortunately, for popular web services like Facebook and Twitter that field billions of requests per day, this would represent petabytes of provenance, making DAP too demanding for widespread use.

To address this problem, we consider a garbage collection technique based on the SQLi defense scenario that was introduced in Section 3.5. In this circumstance, the vast majority of provenance will represent benign web requests. As a result, after the request is whitelisted by the Network Guard mechanism, this information can be discarded. In turn, those requests that were potentially malicious can be written to secondary storage for forensic analysis. By applying this technique, only the provenance of active requests needs be stored in memory. We implemented this garbage collection procedure in our provenance guard, and then repeated the trial. Figure 9 shows that, after the optimization, the provenance store experienced logarithmic, rather than linear, growth. After 6 minutes, the provenance logs grew to just 3.9 MB in size, representing a 75% decrease in storage overhead. While not a total solution to provenance storage overhead, this result indicates that overheads can grow manageably with the size of the web service.

## 6. CASE STUDIES

In this section, we consider several attack scenarios in which DAP can be used to monitor and prevent Internet-based attacks.

### 6.1 Scenario #1: SQL Injection

Through the introduction of the Network Guard component, DAP is able to prevent SQLi-based data exfiltration attacks in real-time by performing provenance-based whitelisting of outbound server responses. As shown in Section 5.3, the Network Guard can authorize (or deny) an outbound message in just a few milliseconds. Here, we procedurally generate the provenance graph of a SQLi attack by using a toy PHP application on Apache 2 that relayed SQL queries to MySQL over POST requests. One example provenance ancestry is shown in Figure 10. Here, we see the provenance of a message derived from SQL data from the customers and orders tables. Any number of the ancestral objects may be an indicator of a data leak. For example, it is unlikely that the web service would explicitly return passwords to the customer. Additionally, a web service would not return a full credit card number to the customer.

Interestingly, the query obfuscations that are commonly associ-



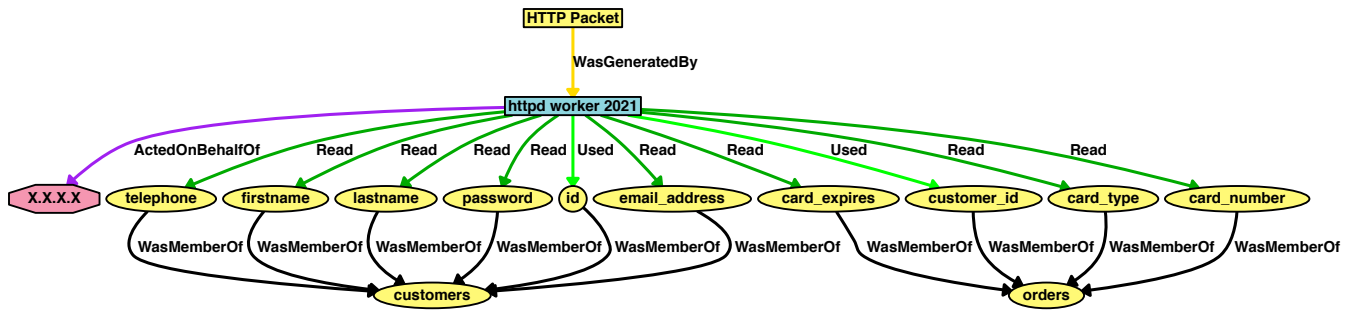


Figure 10: The provenance of a successful SQL injection attack on an eCommerce site launched by remote host X.X.X.X. The attack exfiltrates several valuable data objects from the customers and orders tables.

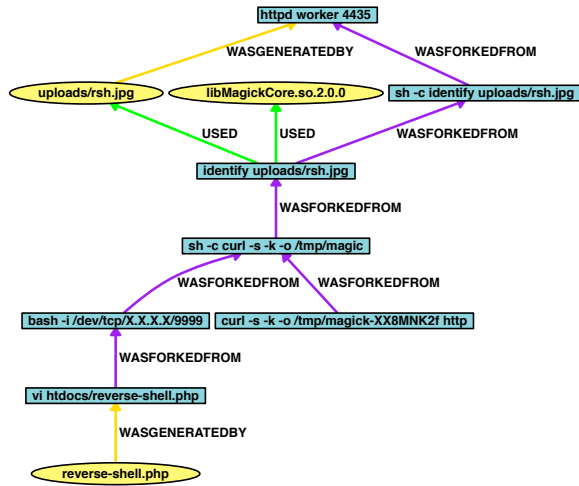


Figure 11: The provenance of an ImageTragick exploit. Network activity has been pruned for clarity. The attacker uploads a malicious image file `rsh.jpg` that opens a remote shell back to the attacker's host. The attacker gains persistence on the machine by placing a reverse shell script in the `htdocs` folder of the server.

ated with SQLi are not present in the provenance graph. This is because such obfuscations are designed to bypass input sanitization checks that are performed by the web application. When a malicious input is able to successfully pass through these checks, the output is a well-formed SQL query. As a result, DAP is optimally positioned to understand the intent of the attacker.

## 6.2 Scenario #2: ImageTragick Exploit

To demonstrate the combined capabilities of DAP and LPM when deployed in tandem, we developed a web application exploit based on the recently discovered vulnerabilities in the ImageMagick image processing library.<sup>10</sup> Our vulnerable web application made use of a PHP script that used the ImageMagick library to test whether an uploaded file was an image. The attack payload was a `jpg` comprised of 4 lines of text including the following command:

```
image over 0,0 0,0
'https://127.0.0.1/x.php?x='bash -i >&
/dev/tcp/X.X.X.X/9999 0>&1''
```

This code is executed by the server when the image is processed by the 'identify' ImageMagick tool, causing a bash shell to be

<sup>10</sup>See <https://imagetrack.com/>

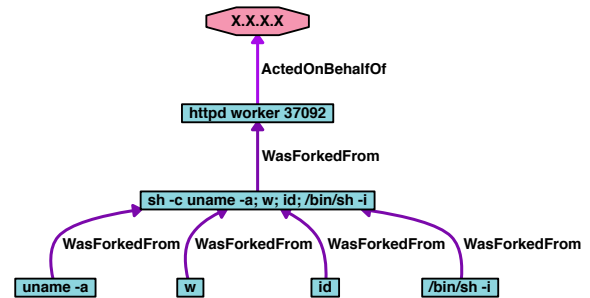


Figure 12: The union of DAP and systems layer provenance can track a remote shell invocation launched by remote host X.X.X.X. The provenance of file and packet manipulations have been pruned for clarity. DAP assists the forensic process by identifying the remote host and unit of work responsible for the exploit.

linked to the attacker's remote host on port 9999. The provenance of these activities on the server side is shown in Figure 11. Without DAP, the operating system provenance for these events would be difficult to interpret. Dependency explosion would make it hard to identify which remote host was able to invoke a shell command. DAP signals the start of a unit of work before the request is handled, which removes from suspicion all sessions that occurred prior to the compromise. Following the compromise, LPM can be securely configured such that the attacker's actions on the system can continue to be monitored [7].

## 6.3 Scenario #3: Reverse Shell Invocation

As a final example, we monitor the attacker's subsequent visits to the web server through invoking the `reverse-shell.php` script. To create realistic attack provenance, we made use of a publicly available php-reverse-shell application.<sup>11</sup> Reverse shells were also an aspect of the Apache 2 Darkleech attack [1]. Figure 12 shows the resulting provenance graph. Here, we can see that an `httpd` worker with pid 37092 is handling a request from remote host X.X.X.X. Unexpectedly, the worker issues a series of conspicuous system commands that collect information regarding the name of the machine (`uname`), the identity of the active user (`id`) which in this case is `daemon`, and the activities of other users that are currently logged in (`w`). The worker then drops to shell. This information would serve as an invaluable explanation as to the attacker's intent once their intrusion had been discovered.

<sup>11</sup>Available at <http://pentestmonkey.net/tools/web-shells/php-reverse-shell>.

## 7. DISCUSSION

*Is DAP a taint tracking system?*

Data provenance and taint analysis are two means of reasoning about information flow. While use of these techniques varies throughout the literature, one means of differentiating between the two is that taint analysis generally provides *what-provenance*, but not *how-provenance*. In a taint analysis solution for SQL Injection, a bit might be flipped in the taint label when a packet contained data from a particular table or column. While this would be sufficient to block the transmission of this packet, it would not provide forensic information about the attack. By providing both *what-* and *how-* provenance, our approach offers a concise explanation of the attack, including identifying the perpetrator.

*Does DAP track individual record accesses and updates?*

The finest granularity that our approach can offer is the granularity of columns; it cannot speak to the specific records that were impacted by a query. This limitation is comparable to the manner in which provenance and audit services track file manipulations at the system call layer; for instance, although the event of a process writing to a file is tracked, the log will not state what was written, or which lines of the file were changed. This choice of granularity represents a design tradeoff between performance and expressivity, and in the design of DAP we have elected to follow convention. We show in Section 3.6 that this granularity can be effective at tracking attacks on web services. Moreover, in order to provide an efficient and rough estimation the impact of a query, we inspect the database engine’s response to the query and measure either its size (in the case of a SELECT) or the number of records affected (in the case of INSERTs and UPDATEs).

*What about NoSQL databases?*

Alternatives to tabular relational databases, including Accumulo, MongoDB, and Cassandra, have become increasingly popular due to their scalability and high performance. While our design focuses on relational databases, we feel that our approach is general enough to apply to these emergent technologies. Rather than extract table column names from queries, DAP for Key-Value stores would instead focus on extracting Key names from queries. It is also worth noting that, in many cases, these systems already provide SQL-like query support.

## 8. RELATED WORK

Our work is part of a growing body of literature that explores the use of provenance to address critical security challenges. Provenance has been employed to detect compromised nodes in data centers [5, 14, 37, 41], explain and prevent data exfiltration [7, 22], and enrich access controls [6, 32].

The notion of provenance tracking originated in literature from the database and scientific workflow communities. Systems such as Chimera offer management of manual provenance annotations [13], but does not perform automatic collection. The Kepler system offers automatic provenance recording for scientific workflows [3], while VisTrails tracks provenance of data visualization procedures [8]. Database management systems such as Trio [38], DBNotes [9], and ORCHESTRA [23] track the provenance of data records as they propagate through the database, and provide custom extensions to SQL so that provenance can be queried. To reduce reliance on custom database engines, the PERM [17] and GProM [4] systems perform automatic query rewriting to annotate result tuples with provenance information.

In spite of the advances made in provenance-aware workflow engines, the systems are “blind” to the rest of the system. They cannot observe information flows beyond the boundaries of their own

operation; critically, this means that they cannot make assertions about whether a given system object (e.g., network packet) was derived from a particular database record. This leaves them of little use when considering the SQL injection and data exfiltration scenarios that motivate our work. In contrast, our system assumes a “black box” database engine, and observes SQL queries and results in order to provide efficient linkability between database operations and other system activity.

### 8.1 Taint Analysis

Like provenance, dynamic taint analysis tracks the propagation of data across a system. Automated instrumentation for taint tracking has been developed for x86 binaries [36], smartphones [11], and databases [20], and dynamic taint analysis in sandbox environments has also been used to secure off-the-shelf applications [42]. Taint tracking systems also suffer from the dependency explosion problem, an effect that can be mitigated, in part, by focusing on short-lived user-facing applications like editors [12]. Provenance can offer a more complete explanation as to *how* an object became tainted. It is also more flexible: taint tracking relies on an immutable policy that requires that data be tagged at runtime, while a provenance-based approach can obtain a result after execution by “replaying” the provenance graph [41], permitting different taints to be considered without re-executing.

### 8.2 Foundational Work

The Nemesis system [10] uses dynamic information flow tracking to prevent authentication bypass attacks. Their system requires manual annotation of the authentication table in the database and then performs taint analysis to track its use in the web application. In contrast, our system can prevent mis-use and exfiltration of any table in the database and does not require web application annotation. Parno et al. built CLAMP [33] to prevent exfiltration from typical web application servers. CLAMP implicitly provides execution partitioning because they create an individual VM for each user’s session. They also provide a database proxy, called the Query Restrictor, which filters queries to the database based on policy. In DAP, we collect the provenance of database accesses at the same location, but deploy policy enforcement in a network guard after the web application has processed the request. This allows us to specify high-level policies that can potentially span multiple database accesses rather than needing to specify policy query by query. Furthermore, we also capture the provenance of the database response, which allows us to enforce policy on the number of records returned.

## 9. CONCLUSION

In spite of a pressing need for ways to explain and mitigate web application vulnerabilities, web services have received little attention as candidates for provenance capabilities. In this work, we presented DAP, a system for creating provenance-aware web applications. Our system can be deployed without requiring any changes to the web application, yet provides rich, concise provenance graphs for web service workflows. We demonstrated DAP’s ability to explain, detect, and prevent SQL injection attacks, and to aid in the tracking of system layer attacks against the web server. In evaluation, we discovered that our system imposes just 5.1 ms overhead on web requests, and introduced a optimization that dramatically reduces the storage burden of provenance capture. Thus, DAP’s non-invasive methodology demonstrates a deployment strategy for provenance not only in web services, but for all complex, heterogeneous application workflows.

## Availability

Our source code and testing framework will be released upon publication.

## 10. REFERENCES

- [1] Darkleech Apache Attacks Intensify. <http://www.darkreading.com/attacks-and-breaches/darkleech-apache-attacks-intensify/d/d-id/1109760?>
- [2] The Second Provenance Challenge. <http://twiki.ipaw.info/bin/view/Challenge/SecondProvenanceChallenge>.
- [3] I. Altintas, O. Barney, and E. Jaeger-Frank. Provenance collection support in the kepler scientific workflow system. In L. Moreau and I. Foster, editors, *Provenance and Annotation of Data*, volume 4145 of *Lecture Notes in Computer Science*, pages 118–132. Springer Berlin Heidelberg, 2006.
- [4] B. Arab, D. Gawlick, V. Radhakrishnan, H. Guo, and B. Glavic. A Generic Provenance Middleware for Database Queries, Updates, and Transactions. June 2014.
- [5] A. Bates, K. Butler, A. Haeberlen, M. Sherr, and W. Zhou. Let SDN Be Your Eyes: Secure Forensics in Data Center Networks. SENT, Feb. 2014.
- [6] A. Bates, B. Mood, M. Valafar, and K. Butler. Towards Secure Provenance-based Access Control in Cloud Environments. In *Proceedings of the 3rd ACM Conference on Data and Application Security and Privacy, CODASPY '13*, pages 277–284, New York, NY, USA, 2013. ACM.
- [7] A. Bates, D. Tian, K. R. Butler, and T. Moyer. Trustworthy Whole-System Provenance for the Linux Kernel. In *Proceedings of 24th USENIX Security Symposium on USENIX Security Symposium*, Aug. 2015.
- [8] S. P. Callahan, J. Freire, E. Santos, C. E. Scheidegger, C. T. Silva, and H. T. Vo. Vistrails: Visualization meets data management. In *Proceedings of the 2006 ACM SIGMOD International Conference on Management of Data, SIGMOD '06*, pages 745–747, New York, NY, USA, 2006. ACM.
- [9] L. Chiticariu, W.-C. Tan, and G. Vijayvargiya. DBNotes: A Post-it System for Relational Databases Based on Provenance. In *Proceedings of the 2005 ACM Special Interest Group on Management of Data Conference, SIGMOD'05*, June 2005.
- [10] M. Dalton, C. Kozyrakis, and N. Zeldovich. Nemesis: Preventing authentication &#38; access control vulnerabilities in web applications. In *Proceedings of the 18th Conference on USENIX Security Symposium, SSYM'09*, pages 267–282, Berkeley, CA, USA, 2009. USENIX Association.
- [11] W. Enck, P. Gilbert, B.-G. Chun, L. P. Cox, J. Jung, P. McDaniel, and A. N. Sheth. TaintDroid: An Information-flow Tracking System for Realtime Privacy Monitoring on Smartphones. In *Proceedings of the 9th USENIX Symposium on Operating Systems Design and Implementation, OSDI'10*, Oct. 2010.
- [12] C. Fleming, P. Peterson, E. Kline, and P. Reiher. Data tethers: Preventing information leakage by enforcing environmental data access policies. In *Communications (ICC), 2012 IEEE International Conference on*, pages 835–840, June 2012.
- [13] I. T. Foster, J.-S. Vöckler, M. Wilde, and Y. Zhao. Chimera: A Virtual Data System for Representing, Querying, and Automating Data Derivation. In *Proceedings of the 14th Conference on Scientific and Statistical Database Management, SSDBM'02*, July 2002.
- [14] A. Gehani, B. Baig, S. Mahmood, D. Tariq, and F. Zaffar. Fine-grained Tracking of Grid Infections. In *Proceedings of the 11th IEEE/ACM International Conference on Grid Computing, GRID'10*, Oct 2010.
- [15] A. Gehani and U. Lindqvist. Bonsai: Balanced Lineage Authentication. In *Proceedings of the 23rd Annual Computer Security Applications Conference, ACSAC'07*, Dec 2007.
- [16] A. Gehani and D. Tariq. SPADE: Support for Provenance Auditing in Distributed Environments. In *Proceedings of the 13th International Middleware Conference, Middleware '12*, Dec 2012.
- [17] B. Glavic and G. Alonso. Perm: Processing Provenance and Data on the Same Data Model Through Query Rewriting. In *Proceedings of the 25th IEEE International Conference on Data Engineering, ICDE '09*, Mar. 2009.
- [18] P. Groth and L. Moreau. Representing Distributed Systems Using the Open Provenance Model. *Future Gener. Comput. Syst.*, 27(6):757–765, June 2011.
- [19] W. G. J. Halfond and A. Orso. Amnesia: Analysis and monitoring for neutralizing sql-injection attacks. In *Proceedings of the 20th IEEE/ACM International Conference on Automated Software Engineering, ASE '05*, 2005.
- [20] W. G. J. Halfond, A. Orso, and P. Manolios. Using positive tainting and syntax-aware evaluation to counter sql injection attacks. In *Proceedings of the 14th ACM SIGSOFT International Symposium on Foundations of Software Engineering, SIGSOFT '06/FSE-14*, pages 175–185, New York, NY, USA, 2006. ACM.
- [21] R. Hasan, R. Sion, and M. Winslett. The Case of the Fake Picasso: Preventing History Forgery with Secure Provenance. In *Proceedings of the 7th USENIX Conference on File and Storage Technologies, FAST'09*, San Francisco, CA, USA, Feb. 2009.
- [22] S. N. Jones, C. R. Strong, D. D. E. Long, and E. L. Miller. Tracking Emigrant Data via Transient Provenance. In *3rd Workshop on the Theory and Practice of Provenance, TAPP'11*, June 2011.
- [23] G. Karvounarakis, T. J. Green, Z. G. Ives, and V. Tannen. Collaborative data sharing via update exchange and provenance. *ACM Trans. Database Syst.*, 38(3):19:1–19:42, Sept. 2013.
- [24] K. H. Lee, X. Zhang, and D. Xu. High Accuracy Attack Provenance via Binary-based Execution Partition. In *Proceedings of the 20th ISOC Network and Distributed System Security Symposium, NDSS'13*, Feb.
- [25] K. H. Lee, X. Zhang, and D. Xu. Loggc: garbage collecting audit log. In *Proceedings of the 2013 ACM SIGSAC conference on Computer &#38; communications security, CCS '13*, pages 1005–1016, New York, NY, USA, 2013. ACM.
- [26] J. Lyle and A. Martin. Trusted Computing and Provenance: Better Together. In *2nd Workshop on the Theory and Practice of Provenance, TAPP'10*, Feb. 2010.
- [27] S. Ma, X. Zhang, and D. Xu. ProTracer: Towards Practical Provenance Tracing by Alternating Between Logging and Tainting. In *Proceedings of the 23rd ISOC Network and Distributed System Security Symposium, NDSS*, Feb. 2016.
- [28] P. Macko and M. Seltzer. A General-purpose Provenance Library. In *4th Workshop on the Theory and Practice of Provenance, TAPP'12*, June 2012.
- [29] K.-K. Muniswamy-Reddy, U. Braun, D. A. Holland, P. Macko, D. Maclean, D. Margo, M. Seltzer, and R. Smogor. Layering in Provenance Systems. In *Proceedings of the 2009 Conference on USENIX Annual Technical Conference, ATC'09*, June 2009.
- [30] K.-K. Muniswamy-Reddy, D. A. Holland, U. Braun, and M. Seltzer. Provenance-aware Storage Systems. In *Proceedings of the Annual Conference on USENIX '06 Annual Technical Conference*, Proceedings of the 2006 Conference on USENIX Annual Technical Conference, June 2006.
- [31] D. Nguyen, J. Park, and R. Sandhu. Dependency Path Patterns As the Foundation of Access Control in Provenance-aware Systems. In *Proceedings of the 4th USENIX Conference on Theory and Practice of Provenance, TAPP'12*, pages 4–4, Berkeley, CA, USA, 2012. USENIX Association.
- [32] J. Park, D. Nguyen, and R. Sandhu. A Provenance-Based Access Control Model. In *Proceedings of the 10th Annual International Conference on Privacy, Security and Trust (PST)*, pages 137–144, 2012.

- [33] B. Parno, J. M. McCune, D. Wendlandt, D. G. Andersen, and A. Perrig. Clamp: Practical prevention of large-scale data leaks. In *Proceedings of the 2009 30th IEEE Symposium on Security and Privacy*, SP '09, pages 154–169, Washington, DC, USA, 2009. IEEE Computer Society.
- [34] D. Pohly, S. McLaughlin, P. McDaniel, and K. Butler. Hi-Fi: Collecting High-Fidelity Whole-System Provenance. In *Proceedings of the 28th Annual Computer Security Applications Conference*, ACSAC'12, Orlando, FL, USA, 2012.
- [35] C. Sar and P. Cao. Lineage File System. <http://crypto.stanford.edu/~cao/lineage.html>.
- [36] P. Saxena, R. Sekar, and V. Puranik. Efficient Fine-grained Binary Instrumentation with Applications to Taint-tracking. In *Proceedings of the 6th Annual IEEE/ACM International Symposium on Code Generation and Optimization*, CGO '08, pages 74–83, New York, NY, USA, 2008. ACM.
- [37] D. Tariq, B. Baig, A. Gehani, S. Mahmood, R. Tahir, A. Aqil, and F. Zaffar. Identifying the Provenance of Correlated Anomalies. In *Proceedings of the 2011 ACM Symposium on Applied Computing*, SAC '11, Mar. 2011.
- [38] J. Widom. Trio: A System for Integrated Management of Data, Accuracy, and Lineage. Technical Report 2004-40, Stanford InfoLab, Aug. 2004.
- [39] World Wide Web Consortium. PROV-Overview: An Overview of the PROV Family of Documents. <http://www.w3.org/TR/prov-overview/>, 2013.
- [40] L. Zhang, A. Persaud, A. Johnson, and Y. Guan. Detection of stepping stone attack under delay and chaff perturbations. In *Performance, Computing, and Communications Conference, 2006. IPCCC 2006. 25th IEEE International*, pages 10 pp.–256, April 2006.
- [41] W. Zhou, Q. Fei, A. Narayan, A. Haeberlen, B. T. Loo, and M. Sherr. Secure Network Provenance. In *Proceedings of the 23rd ACM Symposium on Operating Systems Principles*, SOSP'11, Oct. 2011.
- [42] D. Y. Zhu, J. Jung, D. Song, T. Kohno, and D. Wetherall. TaintEraser: Protecting Sensitive Data Leaks Using Application-level Taint Tracking. *SIGOPS Oper. Syst. Rev.*, 45(1):142–154, Feb. 2011.